

PIGOVIAN TAXES CAN INCREASE PLATFORM COMPETITIVENESS: THE CASE OF ONLINE DISPLAY ADVERTISING

VALERIA STOURM AND ERIC BAX

ABSTRACT. A media platform’s policy on obtrusive ads mediates an effectiveness-nuisance tradeoff. Allowing more obtrusive advertising can increase the effectiveness of ads, so the platform can elicit more short-term revenue from advertisers, but the nuisance to viewers can decrease their engagement over time, which decreases the platform’s opportunity for future revenue. To optimize long-term revenue, a platform can use a combination of advertiser bids and ad impact on user experience to price and allocate ad space.

We study the conditions for advertisers, viewers, and the platform to simultaneously benefit from using ad impact on user experience as a criterion for ad selection and pricing. It is important for advertisers to benefit, because media platforms compete with one another for advertisers. Our results show that platforms with more advertisers competing for ad space are more likely to generate increased profits for themselves and their advertisers by introducing ad impact on user experience as a factor in their auction mechanisms. As a result, doing so can be a successful strategy in competition against other platforms.

Keywords: Platform Competition, Mechanism Design, Advertising, Pigovian Tax

JEL Codes: L1, D4, C7

1. INTRODUCTION

Ads can sometimes be a nuisance to viewers. For example, observing online display ads is associated with lower scores on cognitive tasks and can reduce the number of impressions per viewer (Goldstein et al., 2014). Because obtrusive ads jeopardize a platform’s long-term ability to deliver an engaged audience to advertisers (Goldstein et al., 2014; Chatterjee et al., 2003), platforms such as the Yahoo Ad Exchange have a vested interest in limiting negative ad impact on viewer experience (Anderson and Coate, 2005).

However, distracting or obtrusive ads can be more valuable to some advertisers because they tend to be noticed more. We call this the “effectiveness-nuisance tradeoff.” The value of “higher attentiveness to irritating advertising” has been studied in television commercials: more irritating TV commercials are better remembered and are more likely to be recognized (Aaker and Bruzzone, 1985). Similarly, for internet display advertising, higher attentiveness can make some obtrusive ads more effective at increasing purchase intentions (Goldfarb and Tucker, 2011), click-through rates (Chapman, 2011), and ad recognition and brand attitudes (Yoo and Kim, 2005). Goldfarb and Tucker (2011) estimate that advertisers should be willing to pay 74% more to show these “high-visibility” ads.

Abrams and Schwarz (2008) propose an auction mechanism that uses ad impact on user experience as a criterion for ad selection and a factor in pricing. The mechanism first subtracts from each ad’s bid a charge for the (estimated) reduction in long-term revenue to the platform that would be caused by showing the ad. Then the mechanism conducts a second-price auction based on the adjusted bids. The auction winner is allocated the ad space and pays the auction price plus the charge for ad impact on user experience.

This paper studies the conditions under which switching to the Abrams-Schwarz mechanism from a straightforward second-price mechanism can produce a Pareto improvement – simultaneously improving user experience, platform revenue, and profit for the winning advertiser. Pareto improvement is important for switching to be practical, because platforms compete with one another for advertisers and for users. We find that Pareto improvement

depends on a combination of low correlation between ad effectiveness and ad impact on user experience and having sufficient competition for the platform’s ad space.

Charges for ad impact on user experience can be seen as Pigovian taxes (Pigou, 1932). The idea that imposing a tax can make a marketplace or industry more profitable seems counter-intuitive. If the tax affected all advertisers equally, then it would simply reduce profits for all advertisers. However, the Pigovian tax for ad impact on user experience affects some advertisers more than others, depending on their ads’ impact on user experience. It makes advertisers whose ads seriously harm user experience less competitive by subtracting more from their bids. The reduced competition makes the platform more attractive to the remaining advertisers, who impose little harm on user experience or even improve it. Thus, a Pigovian tax can improve profits for the taxed as a class, if the winners’ gains from reduced competition outweigh the winners’ assessments from the tax.

Pigovian taxes may spur innovation by some advertisers and agencies to create ads that are more attractive to users and also effective. Our analysis does not rely on this to achieve a Pareto improvement in the marketplace. It assumes that advertisers can abandon the platform for competing platforms that do not implement Pigovian taxes. However, our analysis can also be applied to scenarios that include innovation, by capturing the innovation in the model’s joint distributions of advertiser bids and ad impacts on user experience.

The rest of this paper is organized as follows. Section 2 reviews related work. Section 3 reviews the outcomes for advertisers and the platform under auction mechanisms with and without a charge for ad impact on user experience. Section 4 develops general theorems on conditions for switching to an Abrams-Schwarz auction to increase profit for the winning advertiser. Section 5 applies the theorems to a model of the effectiveness-nuisance tradeoff and shows that the charge can be advantageous for advertisers and the platform simultaneously. Section 6 concludes with a discussion of possible directions for future work.

2. RELATED WORK

As mentioned in the introduction, this paper analyzes conditions for an auction mechanism by Abrams and Schwarz (2008) to simultaneously benefit a platform, its users, and advertisers. Abrams and Schwarz show that their mechanism increases the sum of platform and advertiser profits; this paper examines conditions for the mechanism to increase profits for both. The Abrams-Schwarz auction is similar to other auctions, by Kempe and Mahdian (2008) and Wilbur et al. (2013), that impose charges based on how an ad in a stream of ads impacts the audience size and hence the effectiveness of subsequent ads. Goldstein et al. (2014) presents an empirical test to measure the impact of annoying ads on user engagement and a dynamic model to explore how publisher choices to control ad impact on user experience are likely to evolve.

This paper focuses on a platform's ability to switch to a more efficient auction mechanism while competing against other platforms. Early work on platform competition by Katz and Shapiro (1985) shows that if there are strong network effects, then a single platform may be necessary to maximize efficiency. Rochet and Tirole (2003, 2006) define two-sided markets and explore how different forms of competition between platforms affect efficiency and participants' profits. Recent work by Lee (2014) examines conditions for multiple competing platforms to be an equilibrium. Our work focuses on the marginal gains for participants and the platform from switching to a more efficient mechanism. So if the platform is a competitor in equilibrium, our work analyzes sufficient conditions for the platform to maintain or improve its position among competitors by switching.

We show that accounting for ad impact on user experience in auctions for online advertising can increase advertiser profits because it can increase dispersion among bids. Ganuza (2004) and Ganuza and Penalva (2010) use dispersion in a similar way to analyze how much information a seller should release to bidders who have different preferences. In their results as in our ours, dispersion has a more pronounced effect in thinner markets.

The Abrams-Schwarz mechanism is an example of a Pigovian tax (Pigou, 1932). The most common setting for Pigovian taxes is governments seeking to reduce hidden costs (usually

pollution) from industry (Perman et al., 2003; Hildebrand and Plourde, 2001; Oates and Schwab, 1988). If we map publishers to jurisdictions, advertisers to firms in an industry, and users to population, then our results apply to competing jurisdictions with mobile populations, mobile firms, and locally polluting environments (using the terms from Hildebrand and Plourde (2001)). Our results imply that a jurisdiction with sufficient market depth and diversity may, by unilaterally imposing Pigovian taxes, maintain industry profitability and improve the standard of living for its population. In fact, doing so may be an effective strategy for competition against other jurisdictions.

3. MARKETPLACE SETTING AND AUCTION MECHANISMS

This section describes a version of the Abrams-Schwarz (AS) auction for digital display advertising (a single-shot, single-slot auction) and its impact on the outcomes for each class of marketplace participant: viewers, advertisers, and the platform. To understand the effects of the AS auction, we compare it to a standard second-price (SP) auction.

Our descriptions of the AS and SP auction mechanisms use straightforward versions of the mechanisms, simplifying some details found in some actual auctions for online advertising. The bids here are per-impression bids. Some platforms allow advertisers to bid per user response (click, sale, or other action). In these cases, the platforms multiply per-response bids by response probabilities to produce our per-impression bids. So our results apply to those platforms as well. Other simplifications include ignoring reserve prices (Myerson, 1981; Riley and Samuelson, 1981; Ostrovsky and Schwarz, 2009), the fact that winners may be charged a small bid increment above second price, and methods to adjust response-based bids for estimation risk (Bax et al., 2012). These simplifications make it easier to present our analysis and should not have a significant effect on our results. For more information on online advertising marketplaces and mechanisms, refer to Varian (2009) and Edelman et al. (2007). For online markets in general, refer to Levin (2011). For general auction theory, refer to Milgrom (2004) and Nisan et al. (2007).

Begin with the baseline: the SP auction mechanism. SP ranks advertiser offers by their bids (breaking any ties randomly). Let n be the number of bidders, and let b_1, \dots, b_n be their bids. Let w be the index of the maximum bid, and let s be the index of the second-highest bid. SP allocates the advertising opportunity to advertiser w and charges this winner the second price: b_s .

The AS auction builds upon SP by incorporating the hidden cost imposed by an ad's obtrusiveness into the pricing and selection rules. Let z_1, \dots, z_n be measures of the hidden costs imposed on the platform by showing the ad creatives associated with bids b_1, \dots, b_n . We assume that z_i is known to both advertiser i and the platform. This assumption is reasonable because a platform can statistically infer measures of ad hidden costs (Goldstein et al., 2014).

The AS auction mechanism:

- (1) Adjusts offers by subtracting each ad's hidden cost from its bid.
- (2) Holds a second-price auction using the adjusted offers.
- (3) Charges the winner of that auction its hidden cost in addition to its auction price.

In more detail, let \tilde{w} be the index of the highest adjusted offer (the maximum $b_i - z_i$), and let \tilde{s} be the index of the second-highest adjusted offer (the runner up). AS allocates the advertising opportunity to advertiser \tilde{w} and charges this winner the second-highest adjusted offer plus the hidden cost from displaying the winner's ad: $(b_{\tilde{s}} - z_{\tilde{s}}) + z_{\tilde{w}}$. This price is equal to the second-highest bid plus the difference between the hidden costs of the winner and the runner up. Table 1 summarizes the notation for each auction.

Under both auctions, we can assume advertisers bid their private values for the opportunity to advertise: v_1, \dots, v_n . The equilibrium bidding strategies are the same because both auctions are truthful mechanisms. For SP, this is well known, because it is a Vickrey-Clarke-Groves (VCG) auction (Vickrey, 1961; Clarke, 1971; Groves, 1973). The equilibrium bidding strategy for the AS auction is also related to the equilibrium bidding strategy for a VCG auction. Advertiser equilibrium bids for AS are $b_i = b(v_i - z_i) + z_i$, where the bidding function $b(x)$ is the equilibrium bid in a VCG auction for an advertiser with a private value

x (Corollary 2 in Abrams and Schwarz 2008). When one slot is auctioned, $b(x) = x$, so $b_i = v_i$.

Table 2 summarizes the outcomes for each participant class. For simplicity, the information in Table 2 assumes that each auction has at least two nonnegative offers, meaning that $v_i > 0$ for the SP auction and $v_i - z_i > 0$ for the AS auction. In our analysis we assume that advertisers with negative offers ($v_i < 0$) choose not to bid in SP and that advertisers with negative adjusted offers ($v_i - z_i < 0$) choose not to bid in AS.

In both auctions, the winning advertiser's profit is their private value for the opportunity to advertise minus the price paid. Thus, the SP winner earns a profit of $v_w - v_s$ and the AS winner earns $v_{\tilde{w}} - (v_{\tilde{s}} + (z_{\tilde{w}} - z_{\tilde{s}})) = (v_{\tilde{w}} - z_{\tilde{w}}) - (v_{\tilde{s}} - z_{\tilde{s}})$. Notice the similarity: for SP, advertiser profit equals the difference between the top two offers, and for AS, advertiser profit equals the difference between the top two adjusted offers.

Treat private values v_1, \dots, v_n and hidden costs z_1, \dots, z_n as exogenous random variables. Assume private values are drawn i.i.d. over advertisers. Assume there is a joint distribution for v_i and z_i , which may be correlated or dependent on each other. Assume adjusted offers $v_i - z_i$ are drawn i.i.d. over advertisers, based on the joint distribution.

To determine whether SP benefits advertisers, we need to compare the differences between the n th and $(n - 1)$ st order statistics from two different distributions – the distribution of offers, which are the private values, and the distribution of adjusted offers. Furthermore, these distributions are related through the strength of the effectiveness-nuisance tradeoff in the marketplace: if v_i and z_i are highly correlated, then there is a strong effectiveness-nuisance tradeoff because ads with higher offers tend to impose higher hidden costs. The next two sections formalize this comparison and analyze the conditions under which advertisers can benefit from the AS auction.

The outcome for viewers is directly related to the obtrusiveness of the winner's ad, denoted z_w for SP and $z_{\tilde{w}}$ for AS. Notice that AS improves viewer experience by design, because it selects advertising offers based partly on the hidden costs imposed by the impact of ads on viewer experience.

For the platform, total revenue under each auction equals the immediate revenue received from the winning advertiser minus the hidden cost incurred from displaying the winning advertiser's ad creative. (Assume that the hidden costs are assessments of net present value to the platform of ad impact on user experience.) Thus, platform revenue for SP is $v_s - z_w$, and the revenue for AS is $(v_{\tilde{s}} + (z_{\tilde{w}} - z_{\tilde{s}})) - z_{\tilde{w}} = v_{\tilde{s}} - z_{\tilde{s}}$. We analyze conditions for AS to improve platform revenue in Section 4.4. We assume the platform has the same costs of acquiring content and conducting auctions under AS and SP, so an improvement in platform revenue is an improvement in platform profit.

4. ADVERTISER PROFITS: GENERAL THEOREMS

To analyze the conditions under which switching from SP to AS benefits advertisers as a class, this section develops two general theorems that compare advertiser profits under each mechanism. The theorems apply to a general set of joint distributions for private values and hidden costs. They only require that the distribution of adjusted bids (private values minus hidden costs) has a flatter right tail than the distribution of bids (private values). In other words, it only requires that a difference between random variables is more dispersed than one of the random variables.

The theorems use order statistics (see e.g. David and Nagaraja (2003)), relying on the fact that as the number of i.i.d. samples increases, the top two samples are more likely to both occur in the right tail of the distribution. If we sample from a pair of distributions, one for bids and another for adjusted bids, then the distribution with a flatter right tail is likely to have the larger separation between its top two samples when there are many samples. (See Figure 1 for details.)

If ad impacts on viewer experience tend to impose hidden costs on the platform (i.e., $E\{z_i\} > 0$), then the mean of $v_i - z_i$ will be less than the mean of v_i , meaning that AS adjusted bids are on average less than SP bids, as shown in Figure 1. However, AS adjusted bids being less than SP bids does not imply that AS advertiser profit is less than SP advertiser profit, because advertiser profit does not depend directly on the values of the top adjusted bid

and bid. Instead advertiser profit depends on the differences between the top two adjusted bids or bids.

Assume that an advertiser chooses not to bid in SP if $v_i < 0$, and an advertiser chooses not to bid in AS if $v_i - z_i < 0$. Consequently, define SP bids as

$$b_i \equiv \max(v_i, 0) \quad (1)$$

and AS adjusted bids as

$$a_i \equiv \max(v_i - z_i, 0). \quad (2)$$

As shown in Table 2, the profit for the winning advertiser in SP is

$$p_{SP} \equiv b_w - b_s \quad (3)$$

and the profit for the winning advertiser in AS is

$$p_{AS} \equiv a_{\tilde{w}} - a_{\tilde{s}}. \quad (4)$$

Consider the distributions of p_{SP} and p_{AS} , conditional on the number of advertisers. Let f be the pdf and F be the cdf of the SP bids b_i . Similarly, let g be the pdf and G the cdf of the AS adjusted bids a_i . If there are n advertisers, then p_{SP} has the same distribution as

$$F^{-1}(u_1) - F^{-1}(u_2), \quad (5)$$

where u_1 and u_2 are the greatest and second-greatest values among n random variables drawn i.i.d. from $U[0, 1]$. In our case, $u_1 = F(b_w)$ and $u_2 = F(b_s)$. To see why Equation 5 holds, recall the probability integral transform: given a random variable x with a continuous cdf F_x , the random variable $F_x(x)$ is uniformly distributed between zero and one (David and Nagaraja, 2003). Similarly, p_{AS} has the same distribution as

$$G^{-1}(u_1) - G^{-1}(u_2). \quad (6)$$

The following theorems compare p_{AS} to p_{SP} using their distributions from Equations 5 and 6.

Theorem 4.1. *Suppose g (the pdf of adjusted bids) has a lower right tail than f (the pdf of bids) in the sense that*

$$\exists \hat{u} < 1 : \forall u > \hat{u} : g(G^{-1}(u)) < f(F^{-1}(u)). \quad (7)$$

Then

$$Pr \{p_{AS} > p_{SP}\} \rightarrow 1 \text{ as } n \rightarrow \infty. \quad (8)$$

Equivalently,

$$\forall \delta > 0 : \exists \hat{n}(\delta) < \infty : \forall n \geq \hat{n}(\delta) : Pr \{p_{AS} > p_{SP}\} \geq 1 - \delta. \quad (9)$$

Theorem 4.1 maps the propensity for AS profits to be greater than SP profits to \hat{n} , a lower bound on the number of participating advertisers. Theorem 4.2 below extends the likelihood result of Theorem 4.1 to a result comparing the expectations of AS and SP profits for advertisers. It shows that, under the lower right tail condition of Theorem 4.1, if there are a sufficient number of participating advertisers, then p_{AS} is greater in expectation than p_{SP} . This extension shows that AS can offer the winning advertiser greater average profits than SP.

Theorem 4.2. *Suppose g has a lower right tail than f in the sense that*

$$\exists \hat{u} < 1 : \forall u > \hat{u} : g(G^{-1}(u)) < f(F^{-1}(u)), \quad (10)$$

and $F^{-1}(\hat{u}) < \infty$. Then

$$\exists \hat{n}(F, G) < \infty : \forall n \geq \hat{n}(F, G) : E \{p_{AS}\} > E \{p_{SP}\}. \quad (11)$$

Note that both theorems hold under two conditions, which we call the *dispersion* and *market depth* conditions. The dispersion condition states that AS adjusted bids must be more dispersed than SP bids (the flat right tail condition on g). The market depth condition

states that there must be a sufficient number of advertisers ($n > \hat{n}$ from the theorems). The proofs of the theorems are in the appendix.

5. EFFECTIVENESS-NUISANCE TRADEOFF

In this section, we apply the theorems to a model of advertiser private values and hidden costs that incorporates an effectiveness-nuisance tradeoff. In particular, we analyze when the dispersion and market depth conditions hold for this model, so that advertisers as a class benefit from AS. Then we present a result showing when the platform also benefits from AS under the model.

Effectiveness-Nuisance Model. We model the effectiveness-nuisance tradeoff by specifying a joint distribution of advertiser private valuations for an opportunity to advertise and the hidden costs imposed by their ads. This joint distribution determines the distributions of bids and adjusted bids. First, suppose the distribution f of private values that advertisers have for an opportunity to advertise are normal, with mean μ_v and standard deviation σ_v :

$$v_i \sim \mathcal{N}(\mu_v, \sigma_v^2). \quad (12)$$

Suppose ad hidden costs z_i are a convex combination of a fraction c of the private value (cv_i) and $y_i \sim \mathcal{N}(\mu_y, \sigma_y^2)$, a random variable that is normally distributed and independent of v_i :

$$z_i = \theta cv_i + (1 - \theta)y_i. \quad (13)$$

The hidden costs in Equation 13 are a weighted sum between two components, where the weight $\theta \in [0, 1]$ is a measure of the effectiveness-nuisance tradeoff. When $\theta = 0$, z_i is independent of v_i and so there is no effectiveness-nuisance tradeoff. The tradeoff increases with θ , because as θ increases, advertisers with higher valuations tend to impose greater hidden costs. Under this model, the adjusted bids can be written as:

$$v_i - z_i = (1 - c\theta)v_i - (1 - \theta)y_i. \quad (14)$$

So g is a normal distribution as follows:

$$v_i - z_i \sim \mathcal{N}(\mu_a, \sigma_a^2), \quad (15)$$

with mean

$$\mu_a = (1 - c\theta)\mu_v - (1 - \theta)\mu_y, \quad (16)$$

and standard deviation

$$\sigma_a = \sqrt{(1 - c\theta)^2 \sigma_v^2 + (1 - \theta)^2 \sigma_y^2}. \quad (17)$$

In summary, the distribution of bids f is characterized by Equation 12 and the distribution of adjusted bids g is characterized by Equations 15-17. These distributions are linked through θ , the magnitude of the tradeoff. Using this model for the tradeoff, we analyze when the two conditions that are needed for advertisers to benefit from AS hold.

Dispersion Condition. We can determine whether the dispersion condition holds given a magnitude θ of the effectiveness-nuisance tradeoff. First, recall the dispersion condition:

$$\exists \hat{u} < 1 : \forall u > \hat{u} : g(G^{-1}(u)) < f(F^{-1}(u)), \quad (18)$$

where f and F are the pdf and cdf of SP bids and g and G are the pdf and cdf of AS adjusted bids. Also recall that for all $v_i > 0$, SP bids are the private valuations v_i , and for all $v_i - z_i > 0$, AS adjusted bids are the adjusted private valuations $v_i - z_i$.

Let $\hat{u} = \max(F(0), G(0))$. Note that $F(0) = \Pr\{v_i \leq 0\}$ and $G(0) = \Pr\{v_i - z_i \leq 0\}$. Then, for all $u > \hat{u}$, $F^{-1}(u) > 0$ and $G^{-1}(u) > 0$, so $f(F^{-1}(u))$ is equal to the pdf of v_i at $F^{-1}(u)$, and $g(G^{-1}(u))$ is equal to the pdf of $v_i - z_i$ at $G^{-1}(u)$.

The effectiveness-nuisance model specifies that both bids v_i and adjusted bids $v_i - z_i$ are normally distributed, so $f(F^{-1}(u)) = \frac{1}{\sqrt{2\pi}\sigma_v} e^{-\frac{s^2}{2}}$ and $g(G^{-1}(u)) = \frac{1}{\sqrt{2\pi}\sigma_a} e^{-\frac{s^2}{2}}$, where $s = \Phi^{-1}(u)$ ¹. Note that the ratio of pdf tail values at the same cdf value is the inverse ratio

¹ For each value $u > \hat{u}$

$$G^{-1}(u) = \mu_a + s\sigma_a \implies F^{-1}(u) = \mu_v + s\sigma_v. \quad (19)$$

where $s = \Phi^{-1}(u)$, and Φ represents the standard normal CDF. Thus,

$$g(G^{-1}(u)) = \frac{1}{\sqrt{2\pi}\sigma_a} e^{-\frac{s^2}{2}} \text{ and } f(F^{-1}(u)) = \frac{1}{\sqrt{2\pi}\sigma_v} e^{-\frac{s^2}{2}}.$$

of the standard deviations:

$$\forall u > \hat{u} : \frac{g(G^{-1}(u))}{f(F^{-1}(u))} = \frac{\sigma_v}{\sigma_a}. \quad (20)$$

Thus, in the tradeoff model, the dispersion condition holds when $\forall u > \hat{u} : \sigma_a > \sigma_v$.

Now we can determine whether the dispersion condition holds for a given magnitude of the effectiveness-nuisance tradeoff, θ . Substituting for σ_a from Equation 17, $\sigma_a > \sigma_v$ is equivalent to:

$$\sigma_y > \frac{\sqrt{1 - (1 - c\theta)^2}}{1 - \theta} \sigma_v. \quad (21)$$

If $\theta = 0$, the dispersion condition holds because the numerator is zero. In this case, there is no effectiveness-nuisance tradeoff: z_i is independent of v_i . As a result, $v_i - z_i$ is more dispersed than v_i , and AS benefits advertisers. On the other extreme, if $\theta = 1$ then the dispersion condition does not hold because the denominator is zero. In this case, the effectiveness and nuisance are proportional (i.e., $z_i = cv_i$), so for $0 < c < 1$, the adjusted bids ($v_i - z_i$) are less dispersed than the bids (v_i). For $0 < \theta < 1$, the dispersion condition holds if the independent component of z_i (which is $(1 - \theta)y_i$) contributes more variance than the dependent component (which is $c\theta v_i$) subtracts.

Market Depth Required. When the dispersion condition holds, the theorems show that AS benefits advertisers only if a sufficient number of them are interested in participating. How many are needed? The market depth condition can be applied to the tradeoff model to estimate how many advertisers are needed using Z-scores for a standard normal distribution.

First, recall Figure 1 for some intuition. Under our model, the origin 0 may lie within the figure. When a private value v_i is to the left of the origin, the advertiser chooses not to bid in SP. Similarly, when a value $v_i - z_i$ is to the left of the origin, the advertiser chooses not to bid in AS. If charges for the hidden costs make adjusted bids lower than bids on average, then the distribution of $v_i - z_i$ (g in Figure 1) is shifted to the left of the distribution of v_i (f in Figure 1), causing some SP bidders to choose not to bid in AS. To avoid confusion, we will use *participating advertisers* to refer to the n advertisers for whom v_i and z_i values are drawn and *bidders* to refer to those who have nonnegative adjusted bids: $v_i - z_i \geq 0$.

Let's begin with a simple example. From our model of v_i and z_i , suppose that among participating advertisers, at least two have nonnegative $v_i - z_i$ values. Then both bid, and as previously explained in the description of Figure 1, AS is expected to benefit advertisers. However, if only one advertiser has a nonnegative $v_i - z_i$ value, their profit is only that value, not the difference between that value and the runner-up $v_i - z_i$ value. So, in effect, Region A in Figure 1 shrinks to be bordered on the left by zero, which can make it smaller than Region B. In this case advertiser profits are may not be greater under AS, even if the dispersion condition holds.

Thus, for AS to be beneficial in the tradeoff model of v_i and z_i , the number of advertisers n must be approximately large enough to produce a nonnegative second adjusted bid: $a_{\bar{s}} \geq 0$. Because $E\{G(a_{\bar{s}})\} = \frac{n-1}{n+1}$, the number of participating advertisers required to have at least two AS bidders is about n such that

$$G(0) = \frac{n-1}{n+1}. \quad (22)$$

Solve for n and note that $G(0) = Pr\{v_i - z_i \leq 0\} = \Phi(-\frac{\mu_a}{\sigma_a})$, where Φ is the standard normal CDF. Recall that $v_i - z_i \sim \mathcal{N}(\mu_a, \sigma_a)$. Then the lower bound on the number of advertisers needed for AS to benefit advertisers is:

$$n = \frac{2}{\Phi(\frac{\mu_a}{\sigma_a})} - 1. \quad (23)$$

This lower bound varies with θ , the magnitude of the tradeoff, because μ_a and σ_a are functions of θ (Equations 16 and 17). Equation 23 also estimates the minimum number of advertisers required for AS to benefit advertisers as a class. Note that $\frac{\mu_a}{\sigma_a}$ is a Z -score for a standard normal, so we can use standard normal probability tables to find the denominator in Equation 23, producing a lower bound (that depends on θ) on the number of advertisers needed for AS to benefit them as a class.

We illustrate the application of Equation 23 by considering two cases. First case: if the average private value of advertising, v_i , exceeds the average hidden cost, z_i , then

$$\mu_a = E\{v_i - z_i\} > 0, \quad (24)$$

the Z -score is positive, and the probability of bidding is:

$$\Phi\left(-\frac{\mu_a}{\sigma_a}\right) \geq \frac{1}{2}. \quad (25)$$

In this case, it is likely that at least half the participating advertisers will have positive adjusted bids, so only about $n = 3$ participating advertisers are needed for AS to be advantageous.

Second case: if the average hidden cost imposed by ads is greater than the average private value of advertising, then $\mu_a < 0$ so the Z -score is negative. In this case, the relationship between μ_a and σ_a determines approximately how many participants are needed for AS to be advantageous. If $\frac{\mu_a}{\sigma_a} \approx -1$, then each participant's probability of bidding is about the probability of a standard normal being at least one standard deviation above its mean, or about 16%. So about $\frac{1}{6}$ of participants bid in AS, and about 12 participants are needed to have at least two positive adjusted bids. In contrast, if $\frac{\mu_a}{\sigma_a} \approx -3$, then the probability of bidding is about 0.13%, so $n \approx 1500$ advertisers are needed. In between, $\frac{\mu_a}{\sigma_a} \approx -2$ implies that about 90 are needed.

Platform Revenue. As detailed in the introduction, a platform should consider implementing the AS auction when it leads to a Pareto improvement in the marketplace. Thus, we might ask whether the platform can expect to profit from AS when advertisers do. The answer is yes, according to the following theorem, which is proven in the appendix.

Theorem 5.1. *If $c \geq 0$ and $n \geq 3$, AS produces at least as much expected publisher revenue as SP:*

$$\forall c \geq 0, 0 \leq \theta \leq 1, n \geq 3 : E\{v_{\bar{s}} - z_{\bar{s}}\} \geq E\{v_s - z_w\}. \quad (26)$$

Recall that total revenue under AS is $(v_{\tilde{s}} - z_{\tilde{s}}) + z_{\tilde{w}}$ in immediate revenue and $-z_{\tilde{w}}$ in exposure to the hidden cost from displaying ad \tilde{w} , for a total of $v_{\tilde{s}} - z_{\tilde{s}}$. Total revenue under SP is v_s in immediate revenue and $-z_w$ in hidden cost for displaying ad w , for a total of $v_s - z_w$. Intuitively, it is not surprising that AS, by selecting an ad based on total revenue $v_i - z_i$, tends to produce more total revenue than SP, which selects an ad based on only one component of total revenue: v_i .

The theorem requires $c \geq 0$ because our intuition can be misleading when hidden costs z_i are anti-correlated with private values v_i , which occurs when $c < 0$. In this case, it is possible that the platform is already receiving the benefit of selecting an ad with lower hidden costs under SP simply by selecting based on private values, and SP does not compensate advertisers for lowering hidden costs as AS does. For example, suppose $z_i = -\frac{1}{2}v_i$. Then $v_i - z_i = \frac{3}{2}v_i$. So SP and AS have the same winners, and the revenues are $v_s - z_s$ for AS and $v_s - z_w$ for SP. The difference is $-z_s + z_w = \frac{1}{2}(v_s - v_w)$, which is negative.

6. DISCUSSION

This article develops theorems that indicate when a platform can profitably implement the Abrams-Schwarz auction (AS) to improve viewer experience without reducing the expected profits of its advertisers. The theorems show that AS benefits advertisers when dispersion and market depth conditions hold. We find that advertisers tend to benefit more from AS if their ads are less obtrusive than average, if there is a weaker association between ad effectiveness and the hidden costs imposed by their nuisance, and if there are more competing advertisers.

We applied the theorems to a model with an effectiveness-nuisance tradeoff, where advertisers with higher valuations for the opportunity to advertise tend to impose the largest hidden costs on the platform. Applying the theorems to the model yields an estimate of the minimum number of advertisers required to implement AS without reducing advertiser profits. This lower bound can be estimated using a Z -score statistic dependent on advertiser

valuations, ad hidden costs, and the association between them: the effectiveness-nuisance tradeoff.

In addition to showing when the AS auction leads to a Pareto improvement in the marketplace, our results suggest that implementing the AS auction when advertisers are expected to benefit may also be a way for a platform to attract high-quality advertisers. As more platforms implement AS, the incentives for advertisers and creative agencies to develop more ads that are both effective and attractive to viewers will be greater, to the benefit of the entire marketplace.

One nice feature of the theorems is that they impose minimal assumptions on the joint distribution of advertiser valuations and hidden costs, so they can be applied to evaluate scenarios in which this joint distribution is expected to change. For example, in our model for the effectiveness-nuisance tradeoff, the joint distribution is determined by a positive association between advertiser valuations and hidden costs. Suppose that this association is expected to weaken as advertisers invest in developing ads that impose lower hidden costs. With an updated lower estimate of the tradeoff magnitude, the Z -score can still be used to determine how many advertisers are needed for AS to benefit advertisers once the tradeoff decreases in the marketplace.

Another potential application of our theorems is to analyze the implications of implementing AS when competing platforms do so. The market depth condition suggests that when there is an effectiveness-nuisance tradeoff, a platform with a large advertiser base can implement the AS auction to profitably improve viewer experience and expand its lead over competitors. It would be interesting to analyze whether a platform that cannot benefit its advertisers through AS should instead impose a reduced charge for advertiser hidden costs (e.g., subtracting a constant from the charge). This modified AS auction would offer price reductions to advertisers whose ads have the least negative impact on user experience.

REFERENCES

- Aaker, D. A., D. E. Bruzzone. 1985. Causes of irritation in advertising. *Journal of Marketing* **49** 47–57.

- Abrams, Z., M. Schwarz. 2008. Ad auction design and user experience. *Applied Economics Research Bulletin Special Issue I*(Auctions) 98–105.
- Anderson, S. P., S. Coate. 2005. Market provision of broadcasting: A welfare analysis. *Review of Economic Studies* **72**(253) 947–972.
- Bax, E., A. Kuratti, P. McAfee, J. Romero. 2012. Comparing predicted prices in auctions for online advertising. *International Journal of Industrial Organization* **30**(1) 80–88.
- Chapman, M. 2011. What clicks worldwide? rich media can help lift click-through - but only a little. *Adweek (May 31, 2011)* URL <http://www.adweek.com/news/advertising-branding/what-clicks-worldwide-132085>.
- Chatterjee, P., D. L. Hoffman, T. P. Novak. 2003. Modeling the clickstream: Implications for web-based advertising efforts. *Marketing Science* **22**(4) 520–541.
- Clarke, E. H. 1971. Multipart pricing of public goods. *Public Choice* **11**(1) 17–33.
- David, H. A., H. N. Nagaraja. 2003. *Order Statistics*. 3rd ed. Wiley.
- Edelman, B., M. Ostrovsky, M. Schwarz. 2007. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review* **97**(1) 242–259.
- Ganuzza, J.-J. 2004. Ignorance promotes competition: an auction model with endogenous private valuations. *Rand Journal of Economics* 583–598.
- Ganuzza, J.-J., J. S. Penalva. 2010. Signal orderings based on dispersion and the supply of private information in auctions. *Econometrica* **78**(3) 1007–1030.
- Goldfarb, A., C. Tucker. 2011. Online display advertising: Targeting and obtrusiveness. *Marketing Science* **30**(3) 389–404.
- Goldstein, D., S. Suri, P. McAfee, M. Ekstrand-Abueg, F. Diaz. 2014. The economic and cognitive costs of annoying display advertisements. *Journal of Marketing Research* URL <http://journals.ama.org/doi/abs/10.1509/jmr.13.0439>.
- Groves, T. 1973. Incentives in teams. *Econometrica* **41**(4) 617–631.
- Hildebrand, V. A., C. Plourde. 2001. A model of competing jurisdictions with locally polluting environments and mobile populations. *Environmental and Resource Economics* **18** 219–231.
- Katz, M. L., C. Shapiro. 1985. Network externalities, competition, and compatibility. *American Economic Review* **75**(3) 424–440.
- Kempe, D., M. Mahdian. 2008. A cascade model for externalities in sponsored search. *Proceedings of the 4th Workshop on Internet and Network Economics (WINE)* 585–596.
- Lee, R. S. 2014. Competing platforms. *Journal of Economics and Management Strategy* **23**(3) 507–526.
- Levin, Jonathan. 2011. The economics of internet markets. Discussion Papers 10-018, Stanford Institute for Economic Policy Research. URL <http://ideas.repec.org/p/sip/dpaper/10-018.html>.
- Milgrom, P. R. 2004. *Putting Auction Theory to Work*. Cambridge University Press.
- Myerson, R. B. 1981. Optimal auction design. *Mathematics of Operations Research* **6**(1) 58–73.
- Nisan, N., T. Roughgarden, E. Tardos, V. V. Vazirani. 2007. *Algorithmic Game Theory*. 1st ed. Cambridge University Press, New York.
- Oates, W. E., R. M. Schwab. 1988. Economic competition among jurisdictions: Efficiency enhancing or distortion inducing? *Journal of Public Economics* **35** 333–354.

- Ostrovsky, M., M. Schwarz. 2009. Reserve prices in internet advertising auctions: A field study. *Proceedings of the 12th ACM Conference on Electronic Commerce* 59–60.
- Perman, R., Y. Ma, J. McGilvray, M. Common. 2003. *natural Resource and Environmental Economics*. Pearson Education.
- Pigou, A. C. 1932. *Economics of Welfare*. London: Macmillan.
- Riley, J. G., W. F. Samuelson. 1981. Optimal auctions. *American Economic Review* **71**(3) 381–392.
- Rochet, J.-C., J. Tirole. 2003. Platform competition in two-sided markets. *Journal of the European Economic Association* **1**(4) 990–1029.
- Rochet, J.-C., J. Tirole. 2006. Two-sided markets: A progress report. *RAND Journal of Economics* **37**(3) 645–667.
- Varian, H. R. 2009. Online ad auctions. *American Economic Review* **99**(2) 430–434.
- Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* **16**(1) 8–37.
- Wilbur, K. C., L. Xu, D. Kempe. 2013. Correcting audience externalities in television advertising. *Marketing Science* **32**(6) 892–912.
- Yoo, C., K. Kim. 2005. Processing of animation in online banner advertising: The roles of cognitive and emotional responses. *Journal of Interactive Marketing* **19**(4) 18–34.

APPENDIX: PROOFS OF THEOREMS

Proof of Theorem 4.1:

Proof. The random variable u_2 has the distribution $\text{Beta}(n-1, 1)$ (David and Nagaraja, 2003), and

$$\Pr \{ \text{Beta}(n-1, 1) > \hat{u} \} \rightarrow 1 \text{ as } n \rightarrow \infty. \quad (27)$$

(To see this, recall that $E \{ \text{Beta}(n-1, 1) \} = \frac{n-1}{n+1}$ and the Beta distribution is bounded, then apply Markov's inequality.) Because $u_1 > u_2$, this implies

$$\Pr \{ u_1 > \hat{u} \wedge u_2 > \hat{u} \} \rightarrow 1 \text{ as } n \rightarrow \infty. \quad (28)$$

Now we show $u_1, u_2 > \hat{u} \implies p_{AS} > p_{SP}$. Note that

$$\int_{x=F^{-1}(u_2)}^{F^{-1}(u_1)} f(x) dx = u_1 - u_2. \quad (29)$$

So

$$F^{-1}(u_1) - F^{-1}(u_2) = \left\{ d_S \left| \int_{x=F^{-1}(u_2)}^{F^{-1}(u_2)+d_S} f(x) dx = u_1 - u_2 \right. \right\}. \quad (30)$$

In other words, p_{SP} is the distance d_S over which we must integrate f , starting at $F^{-1}(u_2)$, to get $u_1 - u_2$ as the area under the pdf curve. Similarly,

$$G^{-1}(u_1) - G^{-1}(u_2) = \left\{ d_A \left| \int_{x=G^{-1}(u_2)}^{G^{-1}(u_2)+d_A} g(x) dx = u_1 - u_2 \right. \right\}. \quad (31)$$

Because $g(G^{-1}(u)) < f(F^{-1}(u))$ for all $u > \hat{u}$, we must integrate further under g than f to get the same area. So $d_A > d_S$. \square

Proof of Theorem 4.2:

Proof.

$$E \{p_{AS} - p_{SP}\} \quad (32)$$

$$= Pr \{u_2 \leq \hat{u}\} E \{p_{AS} - p_{SP} | u_2 \leq \hat{u}\} + Pr \{u_2 > \hat{u}\} E \{p_{AS} - p_{SP} | u_2 > \hat{u}\} \quad (33)$$

$$\geq Pr \{u_2 \leq \hat{u}\} \min_{u_2 \leq \hat{u}} (p_{AS} - p_{SP}) + Pr \{u_2 > \hat{u}\} E \{p_{AS} - p_{SP} | u_2 > \hat{u}\}. \quad (34)$$

From the proof of Theorem 4.1, the probability in the first term can be made arbitrarily close to zero, and the probability in the second term can be made arbitrarily close to one, by increasing n . The expectation in the second term is positive, because it is a mean over positive values. So we only need to show that the expectation in the first term is bounded below. Note that

$$\min_{u_2 \leq \hat{u}} (p_{AS} - p_{SP}) = \min_{u_2 \leq \hat{u}} \left[(G^{-1}(u_1) - G^{-1}(u_2)) - (F^{-1}(u_1) - F^{-1}(u_2)) \right]. \quad (35)$$

Suppose $u_1 \leq \hat{u}$. The first term in parentheses on the RHS is nonnegative, $F^{-1}(u_2) \geq 0$, and $F^{-1}(u_1) \leq F^{-1}(\hat{u})$, so

$$\min_{u_2 \leq \hat{u}} (p_{AS} - p_{SP}) \geq -F^{-1}(\hat{u}), \quad (36)$$

which is bounded because $F^{-1}(\hat{u}) < \infty$.

Now suppose $u_1 > \hat{u}$. Then

$$(G^{-1}(u_1) - G^{-1}(u_2)) - (F^{-1}(u_1) - F^{-1}(u_2)) \quad (37)$$

$$= [(G^{-1}(u_1) - G^{-1}(\hat{u})) - (F^{-1}(u_1) - F^{-1}(\hat{u}))] + [(G^{-1}(\hat{u}) - G^{-1}(u_2)) - (F^{-1}(\hat{u}) - F^{-1}(u_2))]. \quad (38)$$

The proof of Theorem 4.1 shows that the first term in brackets is nonnegative. The previous case ($u_1 \leq \hat{u}$) shows that the second term is bounded. \square

Proof of Theorem 5.1:

Proof. For expected SP profit:

$$E \{v_s - z_w\} = E \{v_s - z_s\} + E \{z_s - z_w\}. \quad (39)$$

Note that

$$E \{z_s - z_w\} \leq 0, \quad (40)$$

because it is zero if z_i are independent of v_i and negative otherwise. So it remains to show that

$$E \{v_{\bar{s}} - z_{\bar{s}}\} \geq E \{v_s - z_s\}. \quad (41)$$

Recall that

$$v_i - z_i = (1 - c\theta)v_i - (1 - \theta)y_i, \quad (42)$$

where

$$v_i \sim \mathcal{N}(\mu_v, \sigma_v) \quad (43)$$

and

$$y_i \sim \mathcal{N}(\mu_y, \sigma_y). \quad (44)$$

So $v_i - z_i$ has the same distribution as

$$\alpha_i + \beta_i, \quad (45)$$

where

$$\alpha_i \sim \mathcal{N}((1 - c\theta)\mu_v - (1 - \theta)\mu_y, (1 - c\theta)\sigma_v) \quad (46)$$

and

$$\beta_i \sim \mathcal{N}(0, (1 - \theta)\sigma_y). \quad (47)$$

Let s be the index of the runner up α_i and let \tilde{s} be the index of the runner up $\alpha_i + \beta_i$. Then $\alpha_s + \beta_s$ has the same distribution as $v_s - z_s$ and $\alpha_{\tilde{s}} + \beta_{\tilde{s}}$ has the same distribution as $v_{\tilde{s}} - z_{\tilde{s}}$. So we need to show

$$E \{ \alpha_{\tilde{s}} + \beta_{\tilde{s}} \} \geq E \{ \alpha_s + \beta_s \}. \quad (48)$$

Let $n = 3$. Consider any values $\alpha_t \leq \alpha_s \leq \alpha_w$. (We use t for the third place index.) Let $\hat{\alpha}_t$, $\hat{\alpha}_s$, and $\hat{\alpha}_w$ be the reflections of α_t , α_s , and α_w across the mean of the distribution of α_i . For example:

$$\hat{\alpha}_t = \mu_\alpha + (\mu_\alpha - \alpha_t), \quad (49)$$

where

$$\mu_\alpha = (1 - c\theta)\mu_v - (1 - \theta)\mu_y. \quad (50)$$

Because the distribution is normal, the reflections are equally as likely as the original values. Reflection reverses the order, making $\hat{\alpha}_w \leq \hat{\alpha}_s \leq \hat{\alpha}_t$. It preserves distances, making the difference between the winner and runner up in the original values the same as the distance between the runner up and third place in the reflected values and vice versa.

Because β_i has a symmetric zero-mean distribution, β_i and $-\beta_i$ are equally likely. Thus, $\alpha_t + \beta_t$, $\alpha_s + \beta_s$, $\alpha_w + \beta_w$ and $\hat{\alpha}_t - \beta_t$, $\hat{\alpha}_s - \beta_s$, $\hat{\alpha}_w - \beta_w$ are equally likely sets of values for $\alpha_i + \beta_i$. These sets have opposite values of the difference between the runner up among $\alpha_i - \beta_i$ and the $\alpha_i - \beta_i$ value corresponding to the runner up among α_i . Hence

$$E \{ \alpha_{\tilde{s}} + \beta_{\tilde{s}} \} - E \{ \alpha_s + \beta_s \} = 0. \quad (51)$$

Now consider $n > 3$. The runner up α_i is more likely to be to the right of the mean than to the left. Because a normal distribution decreases to the right of the mean, the distance between the runner up and the third place α_i is less than the distance between the winner and runner up, on average:

$$E \{ (\alpha_s - \alpha_t) - (\alpha_w - \alpha_s) \} < 0 \quad (52)$$

Because all β_i have the same variance,

$$E \{ \alpha_{\bar{s}} + \beta_{\bar{s}} \} - E \{ \alpha_s + \beta_s \} > 0. \quad (53)$$

□

FIGURES AND TABLES

Table 1: Notation

Notation	Definition
v_i	advertiser i 's private value for the opportunity to advertise
z_i	advertiser i 's hidden cost
w, s	advertiser indices for the winner and second place (runner up) in SP
\tilde{w}, \tilde{s}	advertiser indices for the winner and second place (runner up) in AS

Table 2: Summary of Outcomes

	SP	AS
Basis for selection	v_i	$v_i - z_i$
Value of slot to winner	v_w	$v_{\tilde{w}}$
Price the winner pays	v_s	$(v_{\tilde{s}} - z_{\tilde{s}}) + z_{\tilde{w}}$
hidden cost incurred by platform	z_w	$z_{\tilde{w}}$
Revenue for platform	$v_s - z_w$	$v_{\tilde{s}} - z_{\tilde{s}}$
Profit for winning advertiser	$v_w - v_s$	$(v_{\tilde{w}} - z_{\tilde{w}}) - (v_{\tilde{s}} - z_{\tilde{s}})$

Advertiser Profits: AS vs. SP




FIGURE 1. Let f be the pdf for bids b_1, \dots, b_n and g be the pdf for adjusted bids a_1, \dots, a_n . AS advertiser profit is $a_{\tilde{w}} - a_{\tilde{s}}$, the difference between the top two samples from g . SP advertiser profit is $b_w - b_s$, the difference between the top two samples from f . (For simplicity, assume the y -axis is to the left of the figure.) Let F and G be the cdfs for f and g . Under the probability integral transform, values $F(b_i)$ and $G(a_i)$ are uniformly distributed over $[0, 1]$. If we draw n independent random variables from a uniform distribution over $[0, 1]$ and order them, the averages from least to greatest are $\frac{1}{n+1}, \dots, \frac{n}{n+1}$. Thus, $E\{F(b_w) - F(b_s)\} = E\{G(a_{\tilde{w}}) - G(a_{\tilde{s}})\} = \frac{n}{n+1} - \frac{n-1}{n+1} = \frac{1}{n+1}$. So the area under the pdf curve g between $a_{\tilde{w}}$ and $a_{\tilde{s}}$ (Region A) must equal the area under f between b_w and b_s (Region B), in expectation. When Regions A and B have the same area, but $g < f$ over corresponding points in the respective regions, $a_{\tilde{w}} - a_{\tilde{s}}$ must be greater than $b_w - b_s$, making AS more profitable than SP.